

Robust Data-Driven Output Feedback Control via Bootstrapped Multiplicative Noise

Benjamin Gravell*

BENJAMIN.GRAVELL@UTDALLAS.EDU

Iman Shames[†]

IMAN.SHAMES@ANU.EDU.AU

Tyler Summers*

TYLER.SUMMERS@UTDALLAS.EDU

*The University of Texas at Dallas

[†]The Australian National University

Editors: R. Firoozi, N. Mehr, E. Yel, R. Antonova, J. Bohg, M. Schwager, M. Kochenderfer

Abstract

We propose a robust data-driven output feedback control algorithm that explicitly incorporates inherent finite-sample model estimate uncertainties into the control design. The algorithm has three components: (1) a subspace identification nominal model estimator; (2) a bootstrap resampling method that quantifies non-asymptotic variance of the nominal model estimate; and (3) a non-conventional robust control design method comprising a *coupled* optimal dynamic output feedback filter and controller with multiplicative noise. A key advantage of the proposed approach is that the system identification and robust control design procedures both use stochastic uncertainty representations, so that the actual inherent statistical estimation uncertainty directly aligns with the uncertainty the robust controller is being designed against. Moreover, the control design method accommodates a highly structured uncertainty representation that can capture uncertainty *shape* more effectively than existing approaches. We show through numerical experiments that the proposed robust data-driven output feedback controller can significantly outperform a certainty equivalent controller on various measures of sample complexity and stability robustness.

Keywords: Robust data-driven control, output feedback, bootstrap, multiplicative noise

1. Introduction

The intersection of data-driven learning and model-based control continues to provide significant research challenges despite its long history and vast research literature. Recent work has focused on non-asymptotic analysis of sample complexity, regret, and robustness, in contrast to a classical focus on asymptotics and stability. Approaches for data-driven control can be broadly divided into two categories: “model-based” (or “indirect”), in which a model for the system dynamics is first learned from data and then used to design a control policy, and “model-free” (or “direct”), in which a control policy is learned directly from data without explicitly learning a model for the system dynamics. Model-based approaches can be further divided into two categories: certainty equivalent, in which uncertainty in the learned model is ignored during control design, and robust, in which uncertainty in the learned model is explicitly accounted for in control design.

Much recent work has considered the full state feedback setting, and some results have very recently been obtained in the partially observed output feedback setting. Finite-sample bounds for system identification from input-state data have been obtained in [Simchowitz et al. \(2018\)](#); [Dean et al. \(2020\)](#) and from input-output data in [Care et al. \(2017\)](#); [Tsiamis and Pappas \(2019\)](#); [Sun et al. \(2020\)](#); [Jedra and Proutiere \(2019\)](#); [Oymak and Ozay \(2021\)](#); [Sarkar et al. \(2021\)](#). Sample

complexity and regret bounds for the Linear Quadratic Gaussian problem are described in [Zheng et al. \(2021\)](#); [Zhang et al. \(2021\)](#) and [Lale et al. \(2020, 2021\)](#); [Simchowitz et al. \(2020\)](#), respectively.

In the partially observed setting, issues around robustness to model uncertainty are much more pronounced than in the full state feedback setting, a fact long known in control theory [Doyle \(1978\)](#). Certainty equivalent approaches that ignore model uncertainty can lead to fragile designs, while existing approaches that incorporate model uncertainty often utilize very coarse uncertainty representations (e.g., spectral norm balls), even when obtaining order optimal statistical sample complexity or regret rates. A good balance between performance and robustness in practice requires carefully constructed and structured uncertainty representations; just as much effort should go into estimating from data the *shape* (not just size) of model uncertainty as the nominal model itself. This becomes especially important as the uncertainty dimension increases: structured uncertainties may have far less volume (in model space) than unstructured ones, thereby enabling superior performance. Developing algorithms with good non-asymptotic performance and robustness properties remains a significant challenge, both in theory and in practice. To address this challenge, [Gravell and Summers \(2020\)](#) proposed a data-driven robust control scheme via bootstrapped multiplicative noise for systems with perfect full state measurements; the present work extends these ideas to the partially observed output feedback setting.

Contributions. The contributions of the present work are as follows:

1. We propose a robust data-driven output feedback control algorithm where the model uncertainty description and robust control design method both use highly structured stochastic uncertainty representations.
2. We present a novel semi-parametric bootstrap algorithm for quantifying structured parametric uncertainty in state space models obtained from subspace identification algorithms using input-output data.
3. We show via numerical experiments that the proposed robust data-driven output feedback controller can significantly outperform a certainty equivalent controller on various measures of sample complexity and stability robustness. We make open-source code implementing the algorithms and experiments freely available.

The algorithm has three components: (1) a subspace identification nominal model estimator; (2) a novel semi-parametric bootstrap resampling method that quantifies non-asymptotic variance of the nominal model estimate; and (3) a non-conventional robust control design method using an optimal linear quadratic coupled estimator-controller with multiplicative noise. This approach provides a natural interface between several highly effective methods from system identification, statistics, and optimal control theory.

1.1. Notation

Symbol	Meaning
$\mathbb{R}^{n \times m}$	Space of real-valued $n \times m$ matrices
$\rho(M)$	Spectral radius (greatest magnitude of an eigenvalue) of a square matrix M
$\ M\ _F$	Frobenius norm (Euclidean norm of the vector of singular values) of a matrix M
$\text{vec}(M)$	Vectorization of matrix M by stacking its columns
$M \succ (\succeq) 0$	Matrix M is positive (semi)definite

An ordered sequence of vectors is denoted in the compact notation $x_{0:T} = [x_0, x_1, \dots, x_T]$.

2. Problem Formulation: Data-Driven Output Feedback Control

We consider data-driven control of the discrete-time linear dynamical system

$$x_{t+1} = Ax_t + Bu_t + w_t, \quad (1)$$

$$y_t = Cx_t + v_t \quad (2)$$

where $x_t \in \mathbb{R}^n$ is the system state, $u_t \in \mathbb{R}^m$ is the control input, $y_t \in \mathbb{R}^p$ is the measured output, and w_t and v_t are i.i.d. process and measurement noises with zero mean and covariance matrices W and V , respectively. The system matrices (A, B, C) and noise covariances (W, V) are grouped into the true model $\mathcal{M} = (A, B, C, W, V)$ which is assumed unknown.¹ Given only on a single training trajectory of finite length T of input-output data $\mathcal{D}_T = (y_{0:T}^{\text{train}}, u_{0:T-1}^{\text{train}})$ generated by the true system (1), (2), a data-driven input-output history-dependent control policy $u_t = \pi(y_{0:t}, u_{0:t-1})$ is to be designed. We assume that the input signal that produced the training trajectory was persistently exciting to avoid identifiability issues (see Definition 5 of [Van Overschee and De Moor \(2012\)](#)).

The performance of an arbitrary policy π is characterized by the infinite-horizon time-averaged linear-quadratic output-input criterion

$$H(\pi) := \lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left[\sum_{t=0}^{T-1} y_t^\top Y y_t + u_t^\top R u_t \right] \quad (3)$$

where $Y \succ 0$ and $R \succ 0$ are penalty matrices, $u_t = \pi(y_{0:t}, u_{0:t-1})$, the initial state x_0 is a random vector with zero mean and identity covariance independent of the noises w_t and v_t , and the expectation is taken with respect to the process and measurement noise sequences and the initial state. Notice that this formulation permits one to choose the output penalty $Y \succ 0$, which can be specified even if the true underlying state x_t and system model are unknown. The output-input performance criterion (3) is equivalent, up to a shift by a problem-dependent constant, to a state-input performance criterion with a penalty matrix $Q = C^\top Y C \succeq 0$, such that

$$J(\pi) := \lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left[\sum_{t=0}^{T-1} x_t^\top Q x_t + u_t^\top R u_t \right] = H(\pi) - \text{Tr}(YV) \quad (4)$$

so that minimization of H is tantamount to minimization of J , which is shifted by a positive constant $\text{Tr}(YV)$ that does not depend on the policy.

We focus on a sequential design pipeline, in which the data is first used to identify a system model $\hat{\mathcal{M}}(\mathcal{D}_T)$ and then an output feedback control policy $\pi_{\hat{\mathcal{M}}(\mathcal{D}_T)}$ is designed based on the identified model; note that the identified model $\hat{\mathcal{M}}(\mathcal{D}_T)$ is more generic and may have alternative or additional structure compared to the true model \mathcal{M} . A *linear dynamic compensator* is a policy which combines a linear state estimator with a linear state estimate feedback in the form

$$\hat{x}_{t+1} = F\hat{x}_t + Ly_t, \quad u_t = K\hat{x}_t. \quad (5)$$

Such a compensator is fully specified by the triple (F, K, L) , and the specification need not depend on the state x_t or system matrices (A, B, C) of the underlying system. The optimal cost is the constant $J^* = \min_{\pi} J(\pi) = J(\pi_{\mathcal{M}})$, which is achieved when the true model \mathcal{M} is known and used in

1. We assume the order n of the underlying system is known; future work will address systems with unknown order.

the canonical linear quadratic Gaussian (LQG) control policy $\pi_{\mathcal{M}}$, a linear dynamic compensator with $F = A + BK - LC$ and gain matrices (K, L) computed (separately) as the solution to two decoupled algebraic Riccati equations, which can be accomplished via several well-known methods such as the dynamic programming techniques of policy iteration and value iteration Bertsekas (2012), convex semidefinite programming Boyd et al. (1994), and specialized direct linear algebraic methods Laub (1979). Therefore, we restrict attention to the class of linear dynamic compensators in (5). Using a compensator (F, K, L) , the closed-loop system dynamics become the autonomous stochastic difference equation

$$\begin{bmatrix} x_{t+1} \\ \hat{x}_{t+1} \end{bmatrix} = \begin{bmatrix} A & BK \\ LC & F \end{bmatrix} \begin{bmatrix} x_t \\ \hat{x}_t \end{bmatrix} + \begin{bmatrix} I & 0 \\ 0 & L \end{bmatrix} \begin{bmatrix} w_t \\ v_t \end{bmatrix} \quad (6)$$

Denote the following augmented closed-loop matrices

$$\Phi = \begin{bmatrix} A & BK \\ LC & F \end{bmatrix}, \quad Q' = \begin{bmatrix} Q & 0 \\ 0 & K^\top RK \end{bmatrix}, \quad W' = \begin{bmatrix} W & 0 \\ 0 & LVL^\top \end{bmatrix}. \quad (7)$$

The stability of the closed-loop system is characterized by the spectrum of the matrix Φ , namely if $\rho(\Phi) < 1$ then the closed-loop system is stable in the sense that the covariance of the augmented state $[x_t \ \hat{x}_t]^\top$ converges to a finite positive definite matrix as $t \rightarrow \infty$. With such stability, the steady-state value matrix P' and the steady-state covariance S' of $[x_t^\top \ \hat{x}_t^\top]^\top$ are found by solving the discrete-time Lyapunov equations

$$P' = \Phi^\top P' \Phi + Q', \quad S' = \Phi S' \Phi^\top + W'. \quad (8)$$

With a slight abuse of notation, the performance criterion (4) can be expressed and computed as

$$J(F, K, L) = \text{Tr}(P'W') = \text{Tr}(S'Q'). \quad (9)$$

Denote the performance criterion and closed-loop system matrix under a linear dynamic compensator $(\hat{F}_T, \hat{K}_T, \hat{L}_T)$ designed with the T -step data record \mathcal{D}_T as $J_T = J(\hat{F}_T, \hat{K}_T, \hat{L}_T)$ from (9) and Φ_T from (7), respectively. The quantity of primary interest is $\frac{J_T}{J_*} \in [1, \infty)$, which represents the normalized infinite-horizon performance at time T . Since the policy is computed based on a model identified from noisy finite data, the ratio $\frac{J_T}{J_*}$ is a random variable. We are interested in its finite sample behavior and finiteness (which relates to stability robustness); in particular, we would like to know not only in how the mean or median scale with the data length T , but also how the upper tails scale. These properties depend on whether and how inherent uncertainty in the identified model is accounted for in the controller design. Certainty equivalent approaches ignore the model uncertainty altogether, which may lead to serious finite sample robustness issues. Here we aim to explicitly incorporate the model uncertainty in the controller design. In particular, we propose a robust data-driven output feedback control algorithm that explicitly accounts for finite-sample model uncertainty in an identified model using a multiplicative noise framework, estimated via the bootstrap.

3. Robust Control via Bootstrapped Multiplicative Noise

Our robust data-driven control algorithm is summarized in Algorithm 1, The algorithm has three main components: (1) a subspace identification nominal model estimator; (2) a bootstrap resampling method that quantifies non-asymptotic variance of the nominal model estimate; and (3) a non-conventional robust control design method using an optimal LQG with multiplicative noise.

Algorithm 1 Robust Data-Driven Output Feedback Control

Input: single trajectory data $\mathcal{D}_T = (y_{0:T}^{\text{train}}, u_{0:T-1}^{\text{train}})$, number of bootstrap resamples N_b , model uncertainty scaling parameter γ , penalty matrices $Y \succ 0, R \succ 0$

- 1: $(\hat{A}_T, \hat{B}_T, \hat{C}_T, \hat{W}_T, \hat{V}_T, \hat{U}_T, \hat{w}_{0:T}, \hat{v}_{0:T}) = \text{SubspaceID}(y_{0:T}, u_{0:T-1})$
 - 2: $(\hat{\Sigma}_{A_T}, \hat{\Sigma}_{B_T}, \hat{\Sigma}_{C_T}) = \text{BootstrapModelCovariance}(y_{0:T}, u_{0:T-1}, \hat{A}_T, \hat{B}_T, \hat{C}_T, \hat{w}_{0:T}, \hat{v}_{0:T}, N_b)$
 - 3: $(\hat{F}_T, \hat{K}_T, \hat{L}_T) = \text{MultiNoiseLQG}(\hat{A}_T, \hat{B}_T, \hat{C}_T, \hat{W}_T, \hat{V}_T, \hat{U}_T, \hat{C}_T^T Y \hat{C}_T, R, \hat{\Sigma}_{A_T}, \hat{\Sigma}_{B_T}, \hat{\Sigma}_{C_T}, \gamma)$
-

3.1. Subspace Identification for Nominal Model Estimation

The first component of the algorithm is a subspace identification algorithm to estimate the unknown system matrices from input-output trajectory data. Subspace identification algorithms have been developed and studied for several decades [Van Overschee and De Moor \(2012\)](#). There are several variations, which all involve constructing block Hankel matrices from the data and estimating certain subspaces via singular value decompositions, from which the system matrices and noise covariances can be retrieved. Any of these can be used within the proposed framework, but for concreteness we use the so-called N4SID algorithm [Van Overschee and De Moor \(1994\)](#). Based on the input-output data $(y_{0:T}^{\text{train}}, u_{0:T-1}^{\text{train}})$, the subspace identification algorithm produces a nominal estimate of the system state space matrices and the process and measurement noise covariances:

$$(\hat{A}_T, \hat{B}_T, \hat{C}_T, \hat{W}_T, \hat{V}_T, \hat{U}_T) = \text{SubspaceID}(y_{0:T}^{\text{train}}, u_{0:T-1}^{\text{train}}). \quad (10)$$

Due to space constraints we refer readers to the literature, e.g. ‘‘Combined Algorithm 1’’ in Chapter 4 of [Van Overschee and De Moor \(2012\)](#), for details of the subspace identification algorithm.

Due to non-uniqueness of state space representations, the system matrices are estimated within a similarity transformation of an underlying unknown representation. Based on the input-output data and the estimated system matrices, subspace algorithms generate *residuals* of the process and measurement noises $\{\hat{w}_\tau\}_{\tau=0}^{t-1}, \{\hat{v}_\tau\}_{\tau=0}^{T-1}$ from which sample average covariance estimates

$$\begin{bmatrix} \hat{W}_T & \hat{U}_T \\ \hat{U}_T^T & \hat{V}_T \end{bmatrix}$$

are produced. Because the estimated system matrices $(\hat{A}_T, \hat{B}_T, \hat{C}_T)$ do not share a state coordinate system with the true system matrices (A, B, C) , even though the true cross-covariance between w_t and v_t is assumed zero, the cross-covariance of the estimates disturbances \hat{w}_t and \hat{v}_t may be non-zero and must be estimated and accounted for in the compensator design.

3.2. Bootstrap Resampling to Quantify Non-Asymptotic Model Uncertainty

There are inevitably errors in model estimates obtained from subspace identification using any finite data record, due to the process and measurement noises. It is difficult to analytically characterize non-asymptotic uncertainty in these estimates. Quantifying uncertainty in subspace identification estimates has been considered in [Viberg et al. \(1991\)](#); [Bauer et al. \(1999\)](#); [Bauer and Jansson \(2000\)](#); [Reynders et al. \(2008\)](#), which focus on asymptotic results. Bootstrapping has been used to quantify non-asymptotic uncertainty in [Bittanti and Lovera \(2000\)](#) for input-output quantities such as frequency response or pole locations. However, to our best knowledge, these uncertainty quantifications have not been used for control design.

To quantify non-asymptotic uncertainty in the model estimate, we propose a novel semi-parametric time series bootstrap resampling procedure. In semi-parametric methods, bootstrap data are simulated from the nominal model with the process and measurement noise sampled i.i.d. with

replacement from residuals calculated with the nominal model [Härdle et al. \(2003\)](#). Dependence in the data is preserved by construction. There are also purely parametric and non-parametric versions of the bootstrap. Generally, the semi- and nonparametric bootstraps are less sensitive to assumptions about the model and the noise distribution, while the semi- and pure parametric bootstraps have better small sample performance when the model is correctly specified. The bootstrap resamples allow for various estimates of finite-sample uncertainty associated with the nominal model; here, we will utilize an estimate of the covariance of the model parameters. For concreteness, a semi-parametric bootstrap with resampled residuals discussed above is summarized in Algorithm 2.

State-space Alignment

Due to non-uniqueness of state space representations, the uncertainty representation should not be obtained directly from a sample covariance of the bootstrap resamples. Instead, for each resample we first find a similarity transformation that minimizes the total squared error to the nominal state space model, and then compute a sample covariance in the transformed coordinates. For brevity, we defer detailed discussion of the state-space alignment to the supplement [Gravell and Summers \(2021\)](#). The coordinate transformation problem is

$$\min_{T \in \mathbb{R}^{n \times n}} d(T) = \psi_A \|T\bar{A} - \hat{A}T\|_F^2 + \psi_B \|T\bar{B} - \hat{B}\|_F^2 + \psi_C \|\bar{C} - \hat{C}T\|_F^2 \quad (11)$$

which can be solved readily and results in an invertible transformation matrix T . Then the transformed system matrices are computed as

$$\tilde{A} = T\bar{A}T^{-1}, \quad \tilde{B} = T\bar{B}, \quad \tilde{C} = \bar{C}T^{-1}, \quad \tilde{W} = T\bar{W}T^\top, \quad \tilde{V} = \bar{V}, \quad \tilde{U} = T\bar{U}. \quad (12)$$

This coordinate alignment is incorporated into the model covariance estimation Algorithm 2.

Algorithm 2 Semi-parametric Bootstrap Model Covariance Estimation

Input: trajectory data $(y_{0:t}, u_{0:t-1})$, nominal model estimate $(\hat{A}_t, \hat{B}_t, \hat{C}_t)$, residuals $\{\hat{w}_\tau\}_{\tau=0}^t, \{\hat{v}_\tau\}_{\tau=0}^t$, number of bootstrap resamples N_b

- 1: $\bar{x}_0 = \hat{x}_0$
- 2: $\bar{u}_{0:t-1} = u_{0:t-1}$
- 3: **for** $k = 1, \dots, N_b$ **do**
- 4: Generate data $\bar{x}_{\tau+1} = \hat{A}_t \bar{x}_\tau + \hat{B}_t \bar{u}_\tau + \tilde{w}_\tau, \bar{y}_\tau = \hat{C}_t \bar{x}_\tau + \tilde{v}_\tau, \tau = 0, \dots, t-1$, where $\tilde{w}_{0:t-1}$ and $\tilde{v}_{0:t-1}$ are i.i.d. resamples with replacement from residuals $\hat{w}_{0:t-1}$ and $\hat{v}_{0:t-1}$
- 5: $(\hat{A}_t^k, \hat{B}_t^k, \hat{C}_t^k, -, -, -, -) = \text{SubspaceID}(\bar{y}_{0:t}, \bar{u}_{0:t-1})$
- 6: $T^* = \arg \min_{T \in \mathbb{R}^{n \times n}} \|T\hat{A}_t - \hat{A}_t^k T\|_F^2 + \|T\hat{B}_t - \hat{B}_t^k\|_F^2 + \|\hat{C}_t - \hat{C}_t^k T\|_F^2$
- 7: $\tilde{A}_t^k = T^* \hat{A}_t T^{*-1}, \quad \tilde{B}_t^k = T^* \hat{B}_t, \quad \tilde{C}_t^k = \hat{C}_t T^{*-1}$
- 8: **end for**

Output: Bootstrap sample covariance $\hat{\Sigma}_{A_t} = \frac{1}{N_b-1} \sum_{k=1}^{N_b} \text{vec}(\tilde{A}_t^k - \hat{A}_t) \text{vec}(\tilde{A}_t^k - \hat{A}_t)^\top$
 Bootstrap sample covariance $\hat{\Sigma}_{B_t} = \frac{1}{N_b-1} \sum_{k=1}^{N_b} \text{vec}(\tilde{B}_t^k - \hat{B}_t) \text{vec}(\tilde{B}_t^k - \hat{B}_t)^\top$
 Bootstrap sample covariance $\hat{\Sigma}_{C_t} = \frac{1}{N_b-1} \sum_{k=1}^{N_b} \text{vec}(\tilde{C}_t^k - \hat{C}_t) \text{vec}(\tilde{C}_t^k - \hat{C}_t)^\top$

3.3. Multiplicative Noise LQG: Combined Controller and State Estimator

The model covariance estimate generated from bootstrap resampling interfaces quite naturally with a variant of the optimal linear quadratic output feedback controller that incorporates multiplicative

noise, which has a long history in control theory but is far less widely known than its additive noise counterpart (Wonham (1967); Bernstein and Greeley (1986); De Koning (1992); Gravell et al. (2019)). Consider the optimal control problem to find an output feedback controller $u_t = \pi(y_{0:t})$ for dynamics perturbed by multiplicative noise

$$\text{minimize}_{\pi \in \Pi} \quad \lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \sum_{t=0}^{T-1} (x_t^\top Q x_t + u_t^\top R u_t), \quad (13)$$

$$\text{subject to} \quad \begin{aligned} x_{t+1} &= (A + \bar{A}_t)x_t + (B + \bar{B}_t)u_t + w_t, \\ y_t &= (C + \bar{C}_t)x_t + v_t \end{aligned} \quad (14)$$

where \bar{A}_t , \bar{B}_t , and \bar{C}_t are i.i.d. zero-mean random matrices with a joint covariance structure over their entries governed by the covariance matrices $\Sigma_A := \mathbb{E}[\text{vec}(\bar{A}) \text{vec}(\bar{A})^\top] \in \mathbb{R}^{n^2 \times n^2}$, $\Sigma_B := \mathbb{E}[\text{vec}(\bar{B}) \text{vec}(\bar{B})^\top] \in \mathbb{R}^{nm \times nm}$, $\Sigma_C := \mathbb{E}[\text{vec}(\bar{C}) \text{vec}(\bar{C})^\top] \in \mathbb{R}^{pn \times pn}$ which quantify uncertainty in the nominal system matrices (A, B, C) . The expectation is taken with respect to all of the basic random quantities in the problem, namely $x_0, \{\bar{A}_t\}, \{\bar{B}_t\}, \{\bar{C}_t\}, \{w_t\}, \{v_t\}$.

Due to the multiplicative noise, the state distribution is non-Gaussian even when all primitive distributions are Gaussian, so the Kalman filter is not necessarily the optimal state estimator. However, the optimal *linear* output feedback controller can be exactly computed, and consists of a multiplicative noise linear dynamic compensator of the form (5). In this case, *there is no separation between estimation and control*, so the optimal controller and estimator gains (K, L) must be jointly computed. Specifically, the optimal gains can be computed by solving the coupled nonlinear matrix equations in symmetric matrix variables $X = (X_1, X_2, X_3, X_4)$ given in the supplementary paper Gravell and Summers (2021) by using a value iteration algorithm, described in De Koning (1992). In the absence of multiplicative noise, they reduce to the familiar separated algebraic Riccati equations for optimal estimation and control. The solutions are denoted

$$(X, K, L) = \text{GDARE}(A, B, C, W, V, U, Q, R, \Sigma_A, \Sigma_B, \Sigma_C) \quad (15)$$

Both the optimal controller and estimator gains depend explicitly on the model uncertainty, as quantified by the variances of the system matrices, as well as the process and measurement noise covariances. This policy is known to provide robustness to uncertainties in the parameters of the nominal model (Bernstein and Greeley (1986)). Furthermore, the uncertainty in the nominal model estimate used in this control design method is richly structured and derived directly from the finite available data.

In the proposed data-driven control algorithm, we simply substitute the estimated nominal model and model covariance matrices obtained from the subspace identification and bootstrap methods into the multiplicative noise compensator design equations. We also introduce a parameter γ which provides a fixed scaling of the model uncertainty. Note that $\gamma = 0$ corresponds to certainty equivalent control, and as γ increases, more weight is placed on uncertainty in the nominal model. For $\gamma \in (0, 1)$, this approach can be interpreted as shrinkage estimation of the model sample covariance matrices towards certainty equivalence Ledoit and Wolf (2004). Existence of a solution to the generalized Riccati equation depends not just on stabilizability and detectability of the nominal system (A, B, C) , but also on the *mean-square stabilizability* via dynamic output feedback of the multiplicative noise system (called *mean-square compensatability* in De Koning (1992)). When the multiplicative noise variances are too large, it may be impossible to stabilize the system in the

mean-square sense. In this case, we scale down the model variances to compute a mean-square stabilizing dynamic output feedback controller; see Algorithm 3. In particular, we verify the system with specified γ is mean-square stabilizable by checking whether the generalized Riccati equation admits a positive semidefinite solution; if not, we find the upper limit $\gamma_{\max} = c_\gamma \gamma$ via bisection (e.g. Burden et al. (1978)) on a scaling $c_\gamma \in [0, 1]$.

Algorithm 3 Multiplicative Noise LQG

Input: Nominal model matrices A, B, C , additive disturbance covariances W, V, U , penalty matrices Q, R , covariances $\Sigma_A, \Sigma_B, \Sigma_C$, scaling γ , bisection tolerance $\epsilon > 0$

1: Find largest $c_\gamma \in [0, 1]$ via bisection such that there exists a feasible solution to (15)

2: $(X, K, L) = \text{GDARE}(A, B, C, W, V, U, Q, R, c_\gamma \gamma \Sigma_A, c_\gamma \gamma \Sigma_B, c_\gamma \gamma \Sigma_C)$

Output: $(A + BK - LC, K, L)$

4. Numerical Experiments

We examined the following 2-state shift register with system, penalty, and noise covariance matrices

$$\left[\begin{array}{c|c|c} A & B & \\ \hline C & & \end{array} \right] = \left[\begin{array}{cc|c} 0 & 1 & 0 \\ 0 & 0 & 1 \\ \hline 1 & -1 & \end{array} \right], \quad \left[\begin{array}{c|c} Q & \\ \hline R & \end{array} \right] = \left[\begin{array}{cc|c} 1 & -1 & \\ -1 & 1 & \\ \hline & & 0.01 \end{array} \right], \quad \left[\begin{array}{c|c} W & \\ \hline V & \end{array} \right] = \left[\begin{array}{cc|c} 0.1 & 0 & \\ 0 & 0.1 & \\ \hline & & 0.1 \end{array} \right]$$

where the output penalty was $Y = 1$ leading to the given value for $Q = C^\top Y C$. The first state stores the previous value of the second state, the second state is determined solely by the control input, and the output is the difference of the two states. This system is based on the one described in Recht (2020), wherein it was shown that the system is extremely sensitive to model identification errors. In particular, despite the open-loop system being perfectly stable with zero eigenvalues, the system under optimal linear quadratic state feedback control is nearly unstable such that any small error in the estimated system matrices produce an unstable closed-loop system. Therefore, this system is likely to see a benefit from the proposed robust control synthesis approach.

The training data \mathcal{D}_T were generated by initializing the state at the origin, applying random controls distributed according to a Gaussian distribution with zero-mean and scaled identity covariance where the scaling was equal to the sum of the largest singular values of W and V (to ensure a sufficiently strong signal-to-noise ratio), and simulating the evolution of the state with the additive process and measurement noise specified by the problem data (W, V) .

For brevity, we abbreviate the control design schemes “certainty-equivalent control” as “CE” and “robust control via multiplicative noise” as “RMN”. To evaluate the performance of RMN relative to CE, we performed Monte Carlo trials to estimate the distribution of several key quantities: infinite-horizon performance, spectral radius of the closed-loop system, model error, and multiplicative noise variances. In each Monte Carlo trial, the actual additive noise disturbances w_t, v_t were drawn independently. The level of additive noise was significant enough that an appreciable number of model estimates remained poor for many timesteps, highlighting the behavior of CE and RMN in the critical high-uncertainty regime. We simulated the system and evaluated quantities for the trajectory lengths $T \in \{20, 40, 80, 160, 320\}$ according to Algorithm 1; all of the trajectory lengths were sufficiently long to ensure the estimates in subspace identification were non-degenerate. We drew $N_s = 100,000$ independent Monte Carlo samples and $N_b = 100$ bootstrap samples at each

time step for uncertainty estimation. We used unity scaling of the multiplicative noise ($\gamma = 1$) and a tolerance of $\epsilon = 0.01$ for bisection to find the largest scaling c_γ of multiplicative noise variance in the multiplicative noise LQG algorithm.

In Figure 1 we plot statistics of performance and spectral radius using both control schemes, while in Figure 2 we plot statistics of the differences between the performance and spectral radius. We are chiefly interested in the expected value and upper quantiles of performance, which correspond to average performance and risk of poor performance. We see that CE leads to both worse average behavior and riskier behavior as reflected by the distribution of the performance. In particular, we see that the performance of RMN is clearly better during times between $T = 20$ and $T = 80$, dropping at the 99th percentile from 2.318 to 1.077 whereas CE suffers 12.091 to 1.089. This can also be explained from the spectral radius, which is larger across all timesteps and statistics, corresponding to a less stable system. In particular, at the beginning between $T = 20$ and $T = 80$ when uncertainty is highest, RMN yields spectral radii at the 99th percentile that drop from 0.903 to 0.843 while CE yields 0.985 to 0.916, nearer to instability and allowing the state to travel far from the origin, resulting in high cost. With increasing T the model estimates improved and uncertainty estimates became sufficiently small that the difference between CE and RMN control was insignificant.

Please see the supplement [Gravell and Summers \(2021\)](#) for additional figures and discussion; in particular, it is shown that the nominal system matrices \hat{A} , \hat{B} , and \hat{C} produced by the subspace identification algorithm approached the true parameters, accompanied by a decrease in the multiplicative noise variances and increase of c_γ to unity. Additionally, the bootstrap distribution of model parameters is shown for a single Monte Carlo trial at $T = 20$, revealing that the bootstrap distribution of models accurately captures the deviation of the nominal model from the true system. Code which realizes the algorithms of this paper and generates the reported results is available from <https://github.com/TSummersLab/robust-adaptive-control-multinoise-output>

5. Conclusions

We proposed a data-driven robust control algorithm that uses the bootstrap to estimate model estimate covariances and a non-conventional multiplicative noise LQG robust output feedback compensator synthesis to explicitly account for model uncertainty. Future work will go towards providing finite-time theoretical performance guarantees using tools from high-dimensional statistics and exploring alternative bootstrap uncertainty quantification schemes and robust control synthesis frameworks based e.g. on linear matrix inequalities and System Level Synthesis.

Acknowledgments

This material is based on work supported by the United States Air Force Office of Scientific Research under award number FA2386-19-1-4073 and the National Science Foundation under award number ECCS-2047040.




Figure 1: Infinite-horizon performance J_T/J^* and closed-loop spectral radius $\rho(\Phi_T)$ vs time for CE and RMN.




Figure 2: Difference between CE and RMN on infinite-horizon performance J_T/J^* and closed-loop spectral radius $\rho(\Phi_T)$ metrics vs time. Differences between metrics using CE and RMN were computed for each Monte Carlo trial individually; statistics of the resulting empirical distribution are shown.

References

- Dietmar Bauer and Magnus Jansson. Analysis of the asymptotic properties of the MOESP type of subspace algorithms. *Automatica*, 36(4):497–509, 2000.
- Dietmar Bauer, Manfred Deistler, and Wolfgang Scherrer. Consistency and asymptotic normality of some subspace algorithms for systems without observed inputs. *Automatica*, 35(7):1243–1254, 1999.
- D Bernstein and S Greeley. Robust controller synthesis using the maximum entropy design equations. *IEEE Transactions on Automatic Control*, 31(4):362–364, 1986.
- Dimitri Bertsekas. *Dynamic programming and optimal control: Volume I*, volume 1. Athena scientific, 2012.
- Sergio Bittanti and Marco Lovera. Bootstrap-based estimates of uncertainty in subspace identification methods. *Automatica*, 36(11):1605–1615, 2000.
- Stephen Boyd, Laurent El Ghaoui, Eric Feron, and Venkataramanan Balakrishnan. *Linear matrix inequalities in system and control theory*. SIAM, 1994.
- Richard L Burden, J Douglas Faires, and Albert C Reynolds. *Numerical analysis*. PWS publishing company, 1978.
- Algo Care, Balázs Cs Csáji, Marco C Campi, and Erik Weyer. Finite-sample system identification: An overview and a new correlation method. *IEEE Control Systems Letters*, 2(1):61–66, 2017.
- W. L. De Koning. Compensability and optimal compensation of systems with white parameters. *IEEE Transactions on Automatic Control*, 37(5):579–588, 1992. doi: 10.1109/9.135491.
- Sarah Dean, Horia Mania, Nikolai Matni, Benjamin Recht, and Stephen Tu. On the sample complexity of the linear quadratic regulator. *Foundations of Computational Mathematics*, 20(4): 633–679, 2020.
- John C Doyle. Guaranteed margins for LQG regulators. *IEEE Transactions on automatic Control*, 23(4):756–757, 1978.
- Benjamin Gravell and Tyler Summers. Robust learning-based control via bootstrapped multiplicative noise. In Alexandre M. Bayen, Ali Jadbabaie, George Pappas, Pablo A. Parrilo, Benjamin Recht, Claire Tomlin, and Melanie Zeilinger, editors, *Proceedings of the 2nd Conference on Learning for Dynamics and Control*, volume 120 of *Proceedings of Machine Learning Research*, pages 599–607. PMLR, 10–11 Jun 2020. URL <https://proceedings.mlr.press/v120/gravell20a.html>.
- Benjamin Gravell and Tyler Summers. Robust data-driven output feedback control via bootstrapped multiplicative noise. -, 2021. URL <https://personal.utdallas.edu/~tyler.summers/papers/RobustDataDrivenOutputFeedback.pdf>.
- Benjamin Gravell, Peyman Mohajerin Esfahani, and Tyler Summers. Learning robust control for LQR systems with multiplicative noise via policy gradient. *CoRR*, 2019. URL <http://arxiv.org/abs/1905.13547>.

- Wolfgang Härdle, Joel Horowitz, and Jens-Peter Kreiss. Bootstrap methods for time series. *International Statistical Review*, 71(2):435–459, 2003.
- Yassir Jedra and Alexandre Proutiere. Sample complexity lower bounds for linear system identification. In *2019 IEEE 58th Conference on Decision and Control (CDC)*, pages 2676–2681. IEEE, 2019.
- Sahin Lale, Kamyar Azizzadenesheli, Babak Hassibi, and Anima Anandkumar. Logarithmic regret bound in partially observable linear dynamical systems. *arXiv preprint arXiv:2003.11227*, 2020.
- Sahin Lale, Kamyar Azizzadenesheli, Babak Hassibi, and Anima Anandkumar. Adaptive control and regret minimization in linear quadratic gaussian (LQG) setting. In *2021 American Control Conference (ACC)*, pages 2517–2522. IEEE, 2021.
- A. Laub. A schur method for solving algebraic Riccati equations. *IEEE Transactions on Automatic Control*, 24(6):913–921, 1979. doi: 10.1109/TAC.1979.1102178.
- Olivier Ledoit and Michael Wolf. Honey, I shrunk the sample covariance matrix. *The Journal of Portfolio Management*, 30(4):110–119, 2004.
- Samet Oymak and Necmiye Ozay. Revisiting Ho-Kalman based system identification: robustness and finite-sample analysis. *IEEE Transactions on Automatic Control*, 2021.
- Ben Recht. Digital witnesses, July 2020. URL <https://www.argmin.net/2020/07/27/discrete-fragility/>.
- Edwin Reynders, Rik Pintelon, and Guido De Roeck. Uncertainty bounds on modal parameters obtained from stochastic subspace identification. *Mechanical systems and signal processing*, 22(4):948–969, 2008.
- Tuhin Sarkar, Alexander Rakhlin, and Munther A Dahleh. Finite time LTI system identification. *J. Mach. Learn. Res.*, 22:26–1, 2021.
- Max Simchowitz, Horia Mania, Stephen Tu, Michael I Jordan, and Benjamin Recht. Learning without mixing: Towards a sharp analysis of linear system identification. In *Conference On Learning Theory*, pages 439–473. PMLR, 2018.
- Max Simchowitz, Karan Singh, and Elad Hazan. Improper learning for non-stochastic control. In *Conference on Learning Theory*, pages 3320–3436. PMLR, 2020.
- Yue Sun, Samet Oymak, and Maryam Fazel. Finite sample system identification: Optimal rates and the role of regularization. In *Learning for Dynamics and Control*, pages 16–25. PMLR, 2020.
- Anastasios Tsiamis and George J Pappas. Finite sample analysis of stochastic system identification. In *IEEE Conference on Decision and Control*, pages 3648–3654. IEEE, 2019.
- Peter Van Overschee and Bart De Moor. N4SID: Subspace algorithms for the identification of combined deterministic-stochastic systems. *Automatica*, 30(1):75–93, 1994. ISSN 0005-1098. doi: [https://doi.org/10.1016/0005-1098\(94\)90230-5](https://doi.org/10.1016/0005-1098(94)90230-5). URL <https://www.sciencedirect.com/science/article/pii/0005109894902305>. Special issue on statistical signal processing and control.

- Peter Van Overschee and BL De Moor. *Subspace identification for linear systems: Theory—Implementation—Applications*. Springer Science & Business Media, 2012.
- Mats Viberg, Björn Ottersten, Bo Wahlberg, and Lennart Ljung. A statistical perspective on state-space modeling using subspace methods. In *[1991] Proceedings of the 30th IEEE Conference on Decision and Control*, pages 1337–1342. IEEE, 1991.
- W Murray Wonham. Optimal stationary control of a linear system with state-dependent noise. *SIAM Journal on Control*, 5(3):486–500, 1967.
- Yifei Zhang, Sourav Kumar Ukil, Ephraim Neimand, Serban Sabau, and Myron E Hohil. Sample complexity of the robust LQG regulator with coprime factors uncertainty. *arXiv preprint arXiv:2109.14164*, 2021.
- Yang Zheng, Luca Furieri, Maryam Kamgarpour, and Na Li. Sample complexity of linear quadratic gaussian (LQG) control for output feedback systems. In *Learning for Dynamics and Control*, pages 559–570. PMLR, 2021.